

A model of **actors** and **grey failures**

*Laura Bocchi, Simon Thompson, Laura Voinea
Julien Lange*

**University of Kent
Royal Holloway, University of London**

Context

- Software execution is affected by **unpredictability**
- Programming practice: handling unpredictability with timeouts, exceptions, supervisors, ...
- Formal models:
 - include reliability primitives e.g., exceptions [Fowler et al. POPL'19], timeouts [Laneve, Zavattaro, FoSSaCS'05][Lopez, Perez, WS-FM'11] ...
 - include link/node failures [Francalanza, Hennessy, CONCUR'05] [Adameir, Peters, Nestmann, FORTE'17]...
 - ...
- Models mostly focus on fail-stop failures

Grey failures

Fail stop

Grey

Byzantine


(a) Permanent Slowdown




(b) Transient Slowdown



(c) Partial Slowdown



(d) Transient Stop




- Multiple “patterns” of failure
- No precise awareness of the state of health of the system
 - A component appears to be working but is experiencing issues
- Differential observation
- Diagnosis is challenging

[Gunavi et al. ACM Transactions on Storage'18]

Contribution

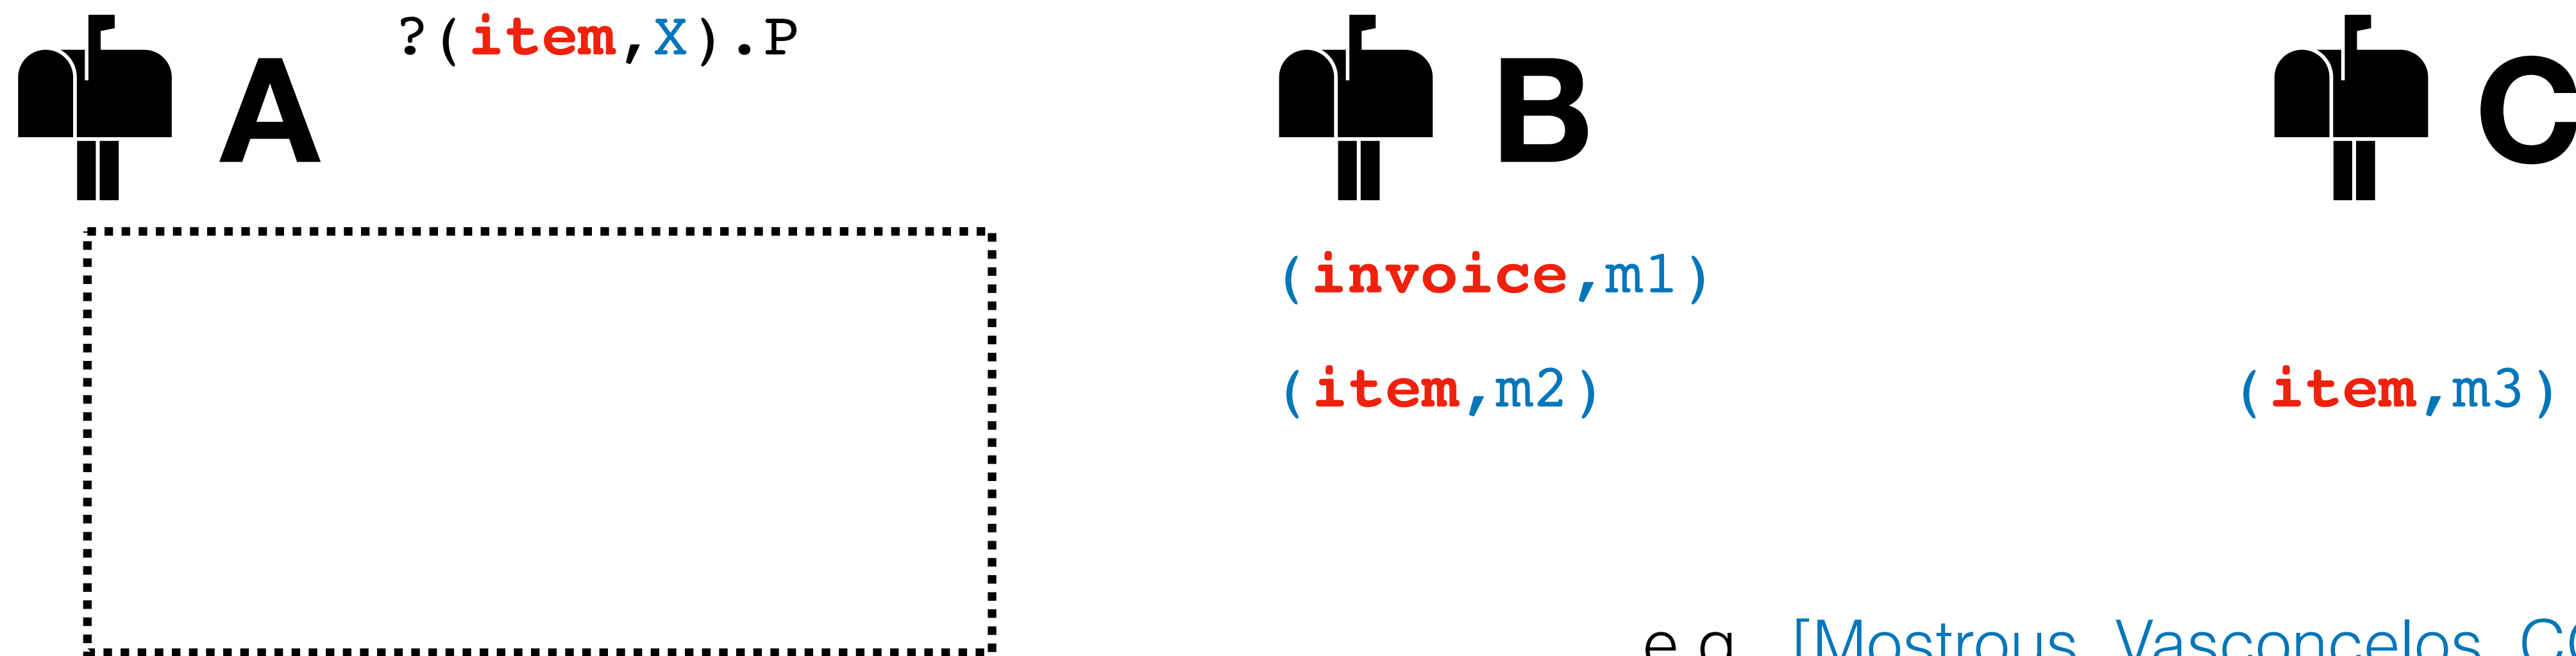
- A formal framework to study the ability of a system to cope with failures
 - a formal model of **actors** and **grey failures**



- a definition of **resilience** and **recoverability**, based on behavioural equivalence

Systems: the main ingredients

- Actor-based systems : we borrow **mailboxes** + **timeouts**
- Mailboxes → expressive communication model e.g., many producers-one consumer
 - asynchrony + pattern-matching + selective receive



e.g., [Mostrous, Vasconcelos, COORDINATION'11]

Systems

$R ::= n[P](M)(t)$
| $R \parallel R$
| 0


$P ::= ?\{p_i \cdot P_i\}_{i \in I}$ **after** P
| $!\{n_i m_i \cdot P_i\}_{i \in I}$
| **sleep** . P
| $\mu t . P$ | t | 0

receive
 pattern1 -> P1;
 ...
 patternN -> PN
after
 n -> P
end

Failures

$\Delta : \mathbb{N} \times \mathcal{N} \cup (\mathcal{N} \times \mathcal{N}) \mapsto \{ \top, \perp, \pm \}$

$\Delta(n)(t) = \begin{cases} \top & \text{if } t = n^2 \ (n \in \mathbb{N}) \\ \perp & \text{otherwise} \end{cases}$



Producer - consumer

instantaneous →

time ↗

average network latency = 1

`p [sleep. !c item. 0](0) || c [?item -> 0 after 3 Cf](0)` time 0

↗ `p [!c item. 0](0) || c [?item -> 0 after 2 Cf](0)` time 1

→ `p 0 || c [?item -> 0 after 2 Cf](0) || sleep.(c,p,item)`

↗ `p 0 || c [?item -> 0 after 1 Cf](0) || (c,p,item)` time 2

→ `p 0 || c [?item -> 0 after 1 Cf](item)`

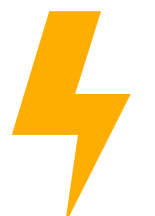
→ `p 0 || c 0`

Producer - consumer with link slowdown

- Reduction on (R, Δ) $\Delta(p, c)(t) = \begin{cases} \pm & \text{if } t \in \{1\} \\ \top & \text{otherwise} \end{cases}$

$R = p [\text{sleep. !c item. 0 }](0) \parallel c [?\text{item} \rightarrow 0 \text{ after 3 Cf }](0)$ time 0

$\rightsquigarrow p [!c \text{ item. 0 }](0) \parallel c [?\text{item} \rightarrow 0 \text{ after 2 Cf }](0)$ time 1

$\rightarrow p 0 \parallel c [?\text{item} \rightarrow 0 \text{ after 2 Cf }](0) \parallel \text{sleep.}(p, c, \text{item})$  (p, c)

$\rightsquigarrow p 0 \parallel c [?\text{item} \rightarrow 0 \text{ after 1 Cf }](0) \parallel \text{sleep.}(p, c, \text{item})$ time 2

$\rightsquigarrow p 0 \parallel c [\text{Cf }](0) \parallel (p, c, \text{item})$ time 3

Behavioural equivalence

Time-abstract weak barbed bisimulation:

$(R_1, \Delta_1) \approx (R_2, \Delta_2)$ implies:

- 1) If $(R_1, \Delta_1) \rightarrow (R'_1, \Delta_1)$ then $(R_2, \Delta_2) \rightarrow^* (R'_2, \Delta_2)$ and $(R'_1, \Delta_1) \approx (R'_2, \Delta_2)$
- 2) If $R_1 \downarrow x$ then $(R_2, \Delta_2) \rightarrow^* (R'_2, \Delta_2)$ and $R'_2 \downarrow x$

Barbs: $n[!\{n_i m_i \cdot P_i\}_{i \in I}](M)(t) \downarrow !n_i m_i$ for all $i \in I$

$(n_1, n_2, m)(t) \downarrow !n_2 m$

$n[?\{p_i \cdot P_i\}_{i \in I} \text{ after } P](M)(t) \downarrow ?n p_i$ for all $i \in I$

Example

$R = \mathbf{p}[\text{sleep} \cdot !\mathbf{c} \text{ item} \cdot 0](0) \parallel \mathbf{c}[\text{?item} \rightarrow 0 \text{ after } 3 \ 0](0) \quad (R, \Delta) \not\approx (R, \top)$

$(R, \top) \rightarrow^* \mathbf{p} 0 \parallel \mathbf{c} 0$

$(R, \Delta) \rightarrow^* \mathbf{p} 0 \parallel (\mathbf{p}, \mathbf{c}, \text{item}) \parallel \mathbf{c}0$

Resilience

- Bisimulation to check the ability of a system to preserve behaviour despite failures
- An initial cursed system (R, Δ) is **resilient** if $(R, \Delta) \approx (R, \top)$
- $\mathbf{p}[\text{sleep}.\!c \text{ item}.0](0) \parallel \mathbf{c}[\text{?item}\rightarrow 0 \text{ after } 3 \ 0](0)$ is **not resilient** (wrt Δ)

$$R_1 = \mathbf{p}[\text{sleep}.\!c \text{ item}.0](0) \parallel \mathbf{c}[\text{?item}\rightarrow 0 \text{ after } 4 \ 0](0) \quad (R_1, \Delta) \approx (R_1, \top)$$

$$R_2 = \mathbf{p}[\text{sleep}.\!c \text{ item}.0](0) \parallel \mathbf{c}[\text{?item}\rightarrow 0 \text{ after } 3 \ (\text{?item}\rightarrow 0 \text{ after } 3 \ 0)](0) \quad (R_2, \Delta) \approx (R_2, \top)$$

$$R_3 = \mathbf{p}[\mu t.\text{sleep}.\!c \text{ item}.t](0) \parallel \mathbf{p}'[\mu t.\text{sleep}.\!c \text{ item}.t](0) \quad (R_3, \Delta) \approx (R_3, \top) \\ \parallel \mathbf{c}[\mu t.\text{?item} \rightarrow 0 \text{ after } 3 \ t](0)$$

Recoverability

- Resilience is too strong to capture e.g., retry strategies

```
p[μt.sleep.!c item.{ok.0, retry.t}](0) ||  
c[μt.?item->!p ok.0 after 3.!p retry.t](0)
```

- An initial cursed system (R, Δ) is ***n*-recoverable** if there exists R' such that $(R, \Delta) \rightarrow^* (R', \Delta)$, $\text{time}(R') = n$, and $(R, T) \approx (R', \Delta)$

Finite thanks to a non-zenoness requirement

Augmentations

- We want to assess resilience/recoverability when **improving** it with recovery strategies

$R = \mathbf{p}[\mathbf{sleep}.\!c \ \mathbf{item}.0](0) \ || \ \mathbf{c}[\?item\rightarrow 0 \ \mathbf{after} \ 3 \ 0](0)$ not resilient wrt Δ

$R_2 = \mathbf{p}[\mathbf{sleep}.\!c \ \mathbf{item}.0](0) \ || \ \mathbf{c}[\?item\rightarrow 0 \ \mathbf{after} \ 3 \ (\?item\rightarrow 0 \ \mathbf{after} \ 3 \ 0)](0)$ resilient wrt Δ

- \bar{R} is an augmentation of R if

(**transparency**) $(\bar{R}, \top) \approx (R, \top)$

(**improvement**) There exists n and Δ s.t. (\bar{R}, Δ) is n -recoverable and (R, Δ) is not

- Moreover, \bar{R} is **preserving** if for all n and Δ ,

(R, Δ) is n -recoverable implies (\bar{R}, Δ) is n -recoverable

Hiding

- Augmentations may add actions e.g., circuit breakers, ... so how about hiding?
- Hiding nodes is not enough (too coarse)

Scoped barbs:

- N is a finite set of elements $!n p$ or $?n p$
- $R \downarrow_N ?n p$ if $R \downarrow ?n p$ and $?n p \notin N$
- $R \downarrow_N ?n m$ if $R \downarrow ?n p$ implies $!n p \not\vdash_{\text{match}} !n m \quad \forall !n p \in N$

A characterisation of failures

- Fundamentals of fault-tolerant distributed computing in asynchronous environments [Gartner 99]
- Simulation relations for fault-tolerance [Demasi, Castro, Maibaum, Aguirre]
- Fault-tolerance:
 - masking (safety+liveness) : $(R, \Delta) \approx (R, T)$
 - fail-safe (safety) : $(R, \Delta) \lesssim (R, T)$
 - non-masking (liveness) : $(R, \Delta) \gtrsim (R, T)$, n-recoverability, ...

Conclusion & future work

- A model for studying **grey failures** in **actor models** (**mailboxes** + **timeouts**)
- Evaluation of recovery strategies reduced to a bisimulation problem
 - resilience, n-recoverability as time-abstract weak barbed bisimulation
 - also augmentation but with existential & universal quantification on n and Δ
- Open questions
 - automated checking and choice of test suites of failures
 - relationships with session types and use in subtyping

Thank you!